Information for your data-protection impact assessment

ZeeDo

Version 1 · Last updated 29 October 2026 · Fingerprint 48beb5020a02

Legal texts are published in English; English binds. This version is kept unchanged. Every version, each with its PDF: https://zeedo.cloud/en/legal/versions/dpia

This page is published by ZeeDo GmbH ("ZeeDo"). Here, "you" and "the Customer" mean the business whose organisation is on ZeeDo, as the Terms of Service define it.

ZeeDo processes personal data for its customers as their processor. In section 8 of the Data Processing Addendum, ZeeDo assists with data-protection impact assessments, taking into account the information available to it, primarily through its documentation. This page is that documentation. It describes the service as it is today, so that you can use it in your own assessment. It is information, not part of the contract and not legal advice: the Data Processing Addendum sets out ZeeDo's obligations. ZeeDo updates this page when the service changes, and the date above shows the version in force.

1. Who decides

You are the controller of the personal data your bots process, and ZeeDo processes it on your behalf. You decide what your bots do, for which purposes and about whom. So you decide whether your use of ZeeDo needs a data-protection impact assessment under Article 35 of the GDPR, and you carry it out. ZeeDo gives you the information on this page and the assistance the Data Processing Addendum describes.

Under the EU AI Act, ZeeDo is the provider of the ZeeDo agent system: it places the system on the market and designs its transparency features. You are its deployer: you decide how, where and for what the system is used, and carry a deployer's duties, including taking measures to support the AI literacy of the people who use the system on your behalf (Article 4) and disclosing AI-generated or manipulated content you publish where Article 50 (4) requires it.

Many uses of AI agents that act towards people will need an assessment. The German supervisory authorities' list of processing operations that always need one names the use of artificial intelligence to process personal data in order to steer the interaction with the people concerned or to assess their personal aspects. The criteria the European Data Protection Board endorsed (WP248 rev.01) include innovative use of new technology, processing on a large scale, sensitive data, data about vulnerable people such as employees, and automated decisions with legal or similarly significant effect; processing that meets two of them will need an assessment in most cases.

Questions to answer for your own use:

- Whom do your bots contact, and on which channels: customers, suppliers, applicants, members of the public?

- Which systems, accounts and websites do your bots read, and what personal data is in them?
- Do your bots handle special categories of data (such as health data, or data revealing religion or trade-union membership), data about criminal offences, or data about children?
- How many people are concerned, how often, and how much data about each?
- Do your bots prepare or take decisions about people that have a legal or similarly significant effect on them?
- Do your bots process your employees' data, or record or observe their work?
- Have the people concerned been told about the processing, as Articles 13 and 14 of the GDPR require?

2. A systematic description of the processing

This section follows Article 35 (7) (a) of the GDPR.

2.1 The service

Each workspace has its own isolated cloud machine, a microVM with its own operating-system kernel. Your bots work on that machine on your instructions. They act through tools, through the apps and accounts you connect, and through a browser and a desktop. For each step, a bot calls an AI model with the context it needs: ZeeDo's model by default, or the model of a provider you choose with your own key (BYOK). Bots keep a memory per workspace, so that they can recall what they learned. ZeeDo records what bots do: each run and its steps, traces of the model calls, screen recordings of the machine, and an audit log of the actions Members and bots take. For security and operation, the servers that run the machines log the commands a bot ran, with their arguments, the paths of the files it wrote (not their contents) and the hosts it reached, for at most 30 days. Bots can send messages to people outside your organisation, such as emails and messages in connected apps.

To make memory and a workspace's tools searchable by meaning, ZeeDo turns memory text, memory searches and tool descriptions into vectors with its own embedding model, and orders search results with its own reranking model. Both run on ZeeDo's servers and keep no copy of the text. An Owner or Admin may instead choose an embedding model of your own OpenAI account for a workspace.

ZeeDo does not use Customer Content to train or fine-tune any AI model, and ZeeDo's model uses only providers that do not train on it.

2.2 The people concerned and their data

The people whose data may be processed are your Members; your employees, customers, suppliers and contacts; anyone whose data appears in the systems or websites you connect your bots to; and the people your bots send messages to.

You determine the categories of personal data through what you ask your bots to do. They may include names, contact details, correspondence, documents, account data at third-party

services, and any data visible on the screens a bot works on. Special categories of data are processed where your bots work with systems or pages that hold them. You decide whether your bots work with such data, and you need your own legal basis and assessment for it. Credentials are stored encrypted in the workspace's vault, and bots use them without reading them.

Annex 1 of the Data Processing Addendum sets out these details as part of the contract.

2.3 Recipients

ZeeDo's sub-processors receive the data they need to run their part of the service. Each is named, with its function and location, on the sub-processor list. Your bots also send data where you direct them: to the apps, accounts and websites they work in and the people they contact. A model provider you use with your own key, an embedding model of your own OpenAI account, and the apps you connect are your own providers, not ZeeDo's sub-processors.

2.4 Locations and transfers

Customer data is stored and processed in the EU (France, Germany and Ireland). Bot machines, ZeeDo's own embedding and reranking models, the control-plane server and the primary file storage run at Scaleway in its Paris region (fr-par), France; a copy of the file storage is kept at OVH in Limburg, Germany; the database (Neon), the workflow service (Temporal Cloud), key management (AWS KMS) and transactional email (Amazon SES) run in Amazon Web Services' Frankfurt region (eu-central-1), Germany; the web application's functions run on Vercel in Frankfurt (fra1), Germany; and traces are kept by Langfuse in its EU region, in Ireland.

ZeeDo's model is the default for every workspace on every plan, Free included, unless a bot or its workspace uses your own key (BYOK) or an Owner or Admin has switched ZeeDo's model off for the workspace. Its requests go to OpenRouter and to the model provider named for each model on the sub-processor list: each model is served by one named provider. Every request is restricted to providers that keep no copy of prompts or responses (zero data retention) and do not train on them, and ZeeDo sets this restriction on each request itself. For an organisation of ZeeDo's EU region, requests go only through OpenRouter's EU endpoint and are processed in the EU; a request that cannot be served there fails and is not sent elsewhere. For an organisation of another region, requests go through the endpoint ZeeDo names for that region.

Two uses may be served from the United States, by a decision model run by Perplexity and reached through OpenRouter, once ZeeDo uses that model for an organisation's region; until then neither is made there, and rules are read by ZeeDo's model as described above. The same restriction applies to both: the provider keeps no copy and does not train on what it receives.

- **Reading rules.** When a person types a message that sets or changes a rule for a bot (with words such as "never", "always" or "from now on"), the typed text of that message and the words of the rules it concerns are sent, to tell whether it is a standing rule; so is the start of a request to a bot whose rules do not all fit in front of it. For this reading, what a person pastes, and anything a bot read, are not sent.
- **Checking a press.** Before a bot presses a button or another control that ZeeDo's own signals

show could pay, order, sign, file, send, publish, delete or change an account, the control is checked, to tell what pressing it would do (4.1). Only labels are sent: the control's label, the labels of its form's fields and of a dialog it answers, the currencies shown, the site's domain and the pattern of the form's address (its path, without the query: a part written only in lower-case letters a to z, hyphens and underscores is kept as written, so it can be a name, such as a user name or a page's name, and every other part is replaced by a placeholder). Field values, the page's text and its links are never sent. Inside a label, every number written in digits (a house number, a postcode, an amount, a date or a reference alike), email addresses, phone numbers, account and card numbers and web addresses written with their scheme (such as https://) or starting with www. are replaced by a placeholder before anything is sent. The other words of a label are sent as written: they may name people, for example the person a message is addressed to, or a street or a place, and an amount written out in words is sent as written. Most presses are not checked by a model at all. Pages of the banking, tax, health and government websites the site registry marks as sensitive, of a site a workspace has marked as sensitive, and of webmail are never sent to the United States: for an organisation of ZeeDo's EU region they are checked in the EU by ZeeDo's model, or not by a model at all. Only the answer and the site's full host name (such as shop.example.com) are kept, never the labels: for at most 30 days, in your organisation's region.

Neither use is made for a workspace on your own key. For an organisation of ZeeDo's EU region, an Owner or Admin can also turn on **Keep every model check in the EU** in Settings: then neither use goes to the United States. Rules are read by ZeeDo's model in the EU as described above, and presses are checked by ZeeDo's model in the same way or, until that check is available, only by ZeeDo's own word lists and the signals in the page. With this setting on, these two checks are made by ZeeDo's model in the EU even for a workspace on your own key, and never by that key's provider.

Members' sign-in data is processed by WorkOS in the United States. The web application's pages and responses pass through Vercel's global edge network, through the edge location nearest the user. Langfuse's group gives support from the United States. OpenRouter is a company in the United States, and several providers that process data in the EU are, or belong to, companies in the United States that may access it from there. Section 6.4 of the Data Processing Addendum names them. Transfers to recipients in, or with access from, a country outside the EEA rest on the Standard Contractual Clauses, as section 6 of the Data Processing Addendum sets out.

2.5 How long data is kept

The main periods for the data your bots process:

Data	Retention
Run history (the step record, not files)	Per plan: Free 7 days, Starter 30, Pro 90, Business 365, Enterprise as agreed. A paid add-on extends it in 90-day steps

Data	Retention
Conversations with bots	On every plan, until the bot or the workspace is deleted. An erasure of a person deletes the lines that name them
Screen recordings	14 days on every plan. Business may extend this to 90 days and Enterprise to 400, never shorten it
Model context: run checkpoints and trace bodies	30 days, on every plan
Workflow history of finished jobs	At most 30 days after the job ends
Operational logs of the servers that run Machines (commands run, paths of files written, hosts reached)	At most 30 days
Memory: raw episodes	At most 90 days
Memory: summaries, facts and relations, procedures, notes people write	Until the Customer deletes them or the workspace is deleted
Audit log and consent records, including approval decisions and standing approvals	400 days on every plan, and at most a month longer, because the log is deleted one calendar month at a time. The plan limits how far back the log can be searched, not how long rows are kept
Backups	Rolling, at most 30 days

Every period is listed in Annex 4 of the Data Processing Addendum and in section 6 of the Privacy Policy. When the service ends, the organisation is read-only for 60 days, its data is deleted on day 90, and backups are purged within 30 days after.

3. Necessity and proportionality: what you control

These settings let you limit what is processed to what your purposes need:

- **Connections and credentials.** You decide which apps and accounts a workspace connects, which credentials its vault holds, and which bots may use them. A bot works only through the connections, credentials and tools its workspace gives it, and on the websites its workspace's web access allows.
- **Web access.** In Settings → Policies & approvals, a workspace either lets bots open any public website that is not blocked (the default), or only the sites you list, each with the methods, paths and request size you allow. You can block further hosts. A site that ZeeDo's own list blocks stays blocked whatever you choose.
- **Approvals.** Every bot starts with each of the six action classes (spend changes, deletions, external messages, merges, publishing and access changes) set to ask a person. You decide who approves. Turning approval off for a class on a bot needs a recent confirmation that it is you and a consent dialog, and you can lock a class at "ask" for a whole workspace.
- **Standing approvals.** A standing approval lets a bot do one narrowly defined thing without

asking each time. It covers one bot, one tool and a condition that only narrows, lasts for a limited time, can be revoked at any time, and ends when the bot's brief, its manager or the connection changes.

- **Memory.** You can see what each bot remembers, stop a bot using an item, or delete an item for good. Raw episodes are deleted after at most 90 days.
- **Screen recordings.** They are kept for 14 days. On Business and Enterprise you may keep them longer; keep them no longer than your purposes need.
- **Model choice.** An Owner or Admin can switch ZeeDo's model off for a workspace, in Settings → Billing & plan, and use your own provider instead (BYOK), under your own contract with that provider.
- **Members and roles.** You decide who belongs to your organisation and in which of the six roles, and so who can see data, instruct bots and approve their actions.
- **Erasure.** An Owner or Admin can request it in Settings → Privacy, or through the public API with an admin-scoped key. An erasure request runs one job across the stores that hold the text of bot memory, chats, runs, cards and approvals, automations, standing approvals and connections, traces, recordings, checkpoints and the files people hand bots. Within 72 hours it permanently deletes what the identifiers given find: in bot memory the matching raw episodes, notes, facts and relations, and the summaries that name the person; the chat lines that name them, each replaced by a marker that a line was erased; the text of the runs, cards and approvals that name them, and those runs' checkpoints and screen recordings; trace bodies; the files handed in those messages and runs; and the text of automations, standing approvals' conditions and connection account names. It then produces a report. Data the identifiers cannot find, such as a face or a name visible only in a screen recording or in a checkpoint of a run whose text never names the person, is deleted when its own period ends (14 days for recordings, 30 days for checkpoints). Deleted data leaves backups within 30 days. Audit and consent records are kept and listed in the report as retained with a reason. The full text of a raw memory episode is stored encrypted and cannot be searched for a name, so an erasure request deletes the full text of every raw episode in every workspace of the organisation; the episodes' records and summaries that do not name the person stay. The job does not search the files bots made or anything else on a Machine's disk, what people wrote to set up a workspace (such as a bot's name and goal), the operational logs of the servers that run Machines, or the history of finished workflows. The report lists the files bots made in the runs concerned, for you to delete. What is on a Machine's disk, and that setup text, stays until a person deletes it or the workspace is deleted; the logs and the workflow history are deleted when their periods end (section 2.5). An export archive made before the request is deleted seven days after it was written. Where a workspace keeps screen recordings longer than 14 days, the recordings the identifiers cannot find are deleted when that longer period ends.
- **Export.** While an organisation is active, and while it is read-only until day 60 of the end-of-service clock, an Owner or Admin can export its data as one archive signed by ZeeDo. The archive is produced within 24 hours and can be downloaded for seven days. Settings → Export can also export only the rows that name one person, to answer that person's request for access. The data export and switching page describes the archive.

4. Risks typical of AI agents, and what ZeeDo's design does about them

For each risk, this section says what ZeeDo's design does and what you do. No technical measure removes these risks completely.

4.1 A bot persuaded by content it reads

Bots read web pages, emails, tickets and tool results written by others, and that content can try to give a bot instructions (prompt injection). No one can yet prevent this completely, and ZeeDo does not promise that a bot resists it.

What ZeeDo's design does. It limits what a misled bot can do. Actions in the six action classes wait for a person's approval unless you turned that off. ZeeDo enforces three controls outside the bot's reach: the approval gate; the payment-form stop and the other refusals of Commitments, which stop a bot at the signs of paying, submitting a regulated filing or defeating a human-verification check that ZeeDo's systems recognise; and the workspace's egress policy. As soon as a run reads content from outside, such as a web page, an email, a tool's answer or another bot's message, ZeeDo marks the run. From then on, standing approvals do not apply to it, a destructive action always waits for a person, any other change it proposes through a tool waits for a person unless a separate check finds it consistent with the bot's brief and the instruction a person gave, and a note it would write to memory is kept only if that check agrees. Before a bot presses a control on a website that ZeeDo's own signals show could matter, ZeeDo checks what pressing it would do (2.4). A press it identifies as paying, ordering, signing or filing is always left to a person. A press it identifies as sending a message, sending a person's details to someone else, publishing, deleting or cancelling something, or changing an account or its security settings follows the bot's settings for that action class, so it may wait for a person's approval. On a marked run, a send, a post or a form with a person's details goes to the separate check, which receives only the kind of action and the site, never the page's labels or content; a deletion or an account change always waits for a person. When a risky press cannot be checked, a person presses the control or approves the press. These checks only add stops. Other actions a bot takes on a website in its browser are not checked this way and do not ask for approval: they are limited by the stops on Commitments, by the rules people set for the bot, and by the websites and accounts you let it use.

What you do. Keep approvals on for actions that matter, grant standing approvals narrowly, connect only what a bot needs, list the sites a bot may open where it needs only a few, do not sign a bot in to an account on a website where it should not act unattended, and do not give a bot a tool or an account that can pay or file without showing the signs ZeeDo's stops recognise.

4.2 Wrong or invented output

Bot output comes from AI models. It can be wrong, incomplete or invented.

What ZeeDo's design does. Work in a regulated category (tax, banking, health, legal) is prepared for your review. A bot is not permitted to file or submit it, and ZeeDo's stops on Commitments apply. Booking, shopping and travel tasks are prepared up to those stops, such as the payment form or the control that places an order, and a person completes them.

What you do. Check output before you rely on it, and have a qualified person check it where it matters.

4.3 A message sent to the wrong person

What ZeeDo's design does. A message a bot sends through a connected message channel, and an email it sends in its browser from a webmail site that ZeeDo recognises to an address outside your organisation, are external messages, one of the six action classes. So every bot starts by asking a person before it sends one, and the person approving sees the recipients and the message, or for a long email its beginning. A press on another website that ZeeDo's check identifies as sending a message or a person's details, for example the send button of a contact form or of a site's own chat, is an external message too and follows the same settings (4.1); the person approving sees the kind of action and the site. Text a bot enters on a website without such a press is not placed in an action class and does not ask for approval.

What you do. Keep approval on for external messages until you trust a bot's work, decide who may approve, and limit the websites and accounts a bot may use where it should not write to people there.

4.4 Collecting more data than needed

What ZeeDo's design does. A bot reaches only the connections, credentials and tools its workspace has, and only the websites its workspace's web access allows: by default, any public website that ZeeDo or you have not blocked.

What you do. Connect only the accounts a bot needs, with the narrowest access the app offers. Delete memory you no longer need, and use the erasure request when a person asks to be forgotten.

4.5 People not knowing they deal with an AI

What ZeeDo's design does. A person a bot contacts is told that they are dealing with an AI agent: ZeeDo's systems add the line — `{bot name}, an AI assistant at {org name}` to the first message of each conversation or thread on each channel, and to the signature of every email a bot sends. The model does not write the line, and no plan or setting switches it off. It is not added to files, to content a person sends, or to a form a bot fills in on a third party's page.

What you do. Tell the people concerned about the processing (Articles 13 and 14 of the GDPR), and disclose AI-generated content you publish where the AI Act requires it.

4.6 Secrets

What ZeeDo's design does. Credentials are stored encrypted in the workspace's vault. Bots use them through the tools you allow, but the model never sees them and no tool returns them to a model. An Owner or Admin can reveal a stored credential after re-authenticating, and each reveal is recorded in the audit log.

What you do. Store only the credentials bots need, give them to the bots that need them, decide

who holds the Owner and Admin roles, and revoke a credential when access should end.

4.7 Access by ZeeDo's people

Everyone ZeeDo authorises to process your personal data is bound to confidentiality, and access to ZeeDo's production systems is limited to the persons ZeeDo authorises to operate the service.

4.8 Personal-data breaches

What ZeeDo does. ZeeDo notifies you of a personal-data breach affecting your data without undue delay, and where feasible within 48 hours of confirming it, as section 9 of the Data Processing Addendum sets out.

What you do. Decide whether to notify your supervisory authority (Article 33 of the GDPR) and the people concerned (Article 34).

4.9 Data leaving the EU

What ZeeDo does. Your data is stored in the EU. These uses take data outside it: the reading of the rules people set for a bot, and the check of what a control a bot is about to press would do, which sends that control's labels only, both of which may be served from the United States unless your organisation, in ZeeDo's EU region, keeps every model check in the EU (2.4); Members' sign-in through WorkOS in the United States; and the delivery of the web application's pages and responses through Vercel's global edge network, through the edge location nearest the user. Data may also be reached from the United States: Langfuse's group gives support from there, and OpenRouter and several of ZeeDo's providers are, or belong to, companies in the United States that may access data from there (section 2.4). These transfers rest on the Standard Contractual Clauses.

What you do. Name these transfers in your records and in the information you give the people concerned.

4.10 Automated decisions and high-risk uses

The Acceptable Use Policy (section 3) does not allow a decision about a person based solely on automated processing, with a legal or similarly significant effect on them, without meaningful review by a person who can change it. ZeeDo is intended for automating a business's office and online work. It is not intended or designed for any use that makes an AI system high-risk under Article 6 (2) and Annex III of the EU AI Act, for example filtering job applications or evaluating candidates or employees, assessing creditworthiness, pricing life or health insurance, deciding on access to education or to essential public services, or identifying people by biometric data. You must not use ZeeDo for such a use. If you do, you act outside ZeeDo's intended purpose and are the provider of that high-risk AI system under Article 25 (1) (c) of the AI Act, with all of a provider's duties.

What you do. Keep a person in every decision about people, and keep your bots to the intended purpose.

4.11 Your employees

If your bots process your staff's data, or recordings show their work, you inform your staff and, where the law requires, involve their representatives (in Germany, for example, a works council) before you start.

5. Technical and organisational measures

Annex 2 of the Data Processing Addendum lists the technical and organisational measures, and section 11.2 of the Data Processing Addendum says how ZeeDo tests its own security. ZeeDo holds no third-party certification or attestation report.

6. What ZeeDo cannot assess for you

Only you can assess:

- your purposes, and whether the processing is necessary and proportionate for them;
- your legal bases for each processing operation, including any special categories of data;
- what the people concerned can reasonably expect;
- the content of the systems, accounts and websites you connect your bots to;
- your staff, their representatives and any works-council agreement;
- the risk that remains after your measures, and whether you must consult your supervisory authority (Article 36 of the GDPR).

7. Contact

Questions about this page or about data protection at ZeeDo: privacy@zeedo.cloud. Assistance beyond this page and ZeeDo's documentation is provided under section 8 of the Data Processing Addendum.